REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-03-

0194

1. REPORT DATE *(DD-MM-YYYY)* 04-09-2002	**2. REPORT TYPE** ~~Annual Progress Report~~	**3. DATES COVERED** *(From - To)* Jan 1 2000-November 30 2002
4. TITLE AND SUBTITLE Attentive Sensing	Final Techn Report	**5a. CONTRACT NUMBER**
		5b. GRANT NUMBER F49620-00-1-0124
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Darryl Morrell		**5d. PROJECT NUMBER**
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Electrical Engineering Arizona State University Tempe AZ 85287-5706		**8. PERFORMING ORGANIZATION REPORT NUMBER** TRC-DM-0301
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Dr. Jon Sjogren Air Force Office of Scientific Research Room 713 4015 Wilson Blvd Arlington VA 22203-1954		**10. SPONSOR/MONITOR'S ACRONYM(S)** AFOSR
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Distribution Unlimited

Please accept as Final Tech & retire the folder jas

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This report is the final progress report for a project whose goal is to examine a class of problems related to the effective allocation of sensing resources in situations involving tunable sensors, configurable sensor suites, or constraints in communication bandwidth and processing resources. We characterized the performance of optimal and suboptimal sensor allocation strategies for target location and detection problems. We have also implemented and characterized the performance of a target tracking algorithm that uses a configurable foveal sensor.

20030612 032

15. SUBJECT TERMS

Sensor Configuration, Optimal Bayesian Control, Heuristic Control

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Darryl Morrell
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL	21	**19b. TELEPHONE NUMBER** *(include area code)* 480-965-2045

TRC Report No. TRC-DM-0301

Attentive Sensing
Final Report
January 1, 2000 through November 30, 2002

Darryl Morrell-PI (January 1 2001-November 30, 2002)
Doug Cochran-PI (January 1 2000-December 31 2000)

Department of Electrical Engineering
Arizona State University
Tempe AZ 85287-5706

AFOSR Grant F49620-00-1-0124

1 Executive Summary

The overall objective of the Attentive Sensing project was to examine a class of problems involving the effective allocation of sensing resources in situations involving tunable sensors, configurable sensor suites, or constraints in communication bandwidth and processing resources. In this executive summary, we briefly summarize the technical accomplishments in this project from January 1, 2000 to November 30. We also describe the personnel supported by this project, the publications resulting from the project, and technical interactions arising from this project.

1.1 Technical Accomplishments

The technical work done in this project was organized into two broad areas. The first was attentive detection and localization of a target using a configurable sensor; this area dealt primarily with problems in which the unknowns and sensor configuration variables were discrete. The second area was attentive estimation of the state of a dynamic system; this area involved unknowns and sensor configuration variables that were continuous. We briefly describe the progress in each of these two areas.

1.1.1 Detection and Localization of Targets

Our objective in this area was to formulate and evaluate optimal and good suboptimal sensor configuration strategies for target detection and localization using a configurable sensor. We used the following problem formulation to investigate sensor configuration strategies: prior information about the target presence and location is represented by a probability distribution; the sensor performance in a given configuration is characterized in terms of probabilities of detection and false alarm; and a loss function is used to represent the utility of potential decisions. In a typical problem, a target, located in one of M cells, is detected and/or localized by repeated use of a sensor to interrogate cells. The sensor makes both missed detection and false alarm errors. Both single mode and dual mode sensors were investigated. Sensor configuration rules were derived with respect to two optimality criteria:

1. Minimize the probability of error when the sensor is applied a fixed number of times (*i.e.* over a fixed time horizon).

2. Minimize the number of sensor uses needed to achieve a given probability of error.

We implemented Matlab and C++ programs to compute the performance of optimal and sub-optimal sensor scheduling rules for both optimality criteria.

For Criterion 1, we investigated the multi-use optimal configuration rule as well as several heuristic sensor scheduling algorithms; the heuristics were all myopic (*i.e.* they determine only the next sensor configuration, not the sequence of sensor configurations). The computational cost of the myopic rule increases linearly with the length of the time horizon, while the cost of the multi-step rule increases exponentially with the length of the time horizon. Thus, for computational reasons, the multi-use optimal rule was investigated only for small number of cells and for short time horizons. The evaluated heuristics include:

1. Interrogate the most probable cell.

2. Interrogate the second most probable cell.

3. Interrogate the myopic optimal cell (*i.e.* the cell which leads to the lowest probability of error after a single interrogation).

We found that for this problem, the heuristic of choosing the myopic optimal cell performs better than the other two heuristics, and that this heuristic gives near optimal performance (compared to the multi-step optimal rule) for both single and dual mode sensors.

For Criterion 2, we developed a method of computing an approximately optimal sensor configuration rule to minimize the number of sensor uses needed to achieve a given probability of error; this method is based on the discretization of the target probability simplex to a finite grid, and is limited for computational reasons to small numbers of cells. We also investigated the three heuristics described above and a two stage heuristic that first detects and then localizes the target. We compared the performance of the most probable heuristic with that of the optimal rule to minimize the number of sensor uses, and found that the performance of the most probable heuristic was good but clearly suboptimal. We also found that the most probable heuristic gave the best performance of the four heuristics evaluated.

1.1.2 Attentive Estimation of the State of a Stochastic System

Our objective in this areas was to find and evaluate optimal and good suboptimal configuration strategies for the problem of estimating the state of a stochastic system from measurements obtained from a configurable sensor. We used the problem of target tracking with a non-linear "foveal sensor" for this investigation.

A foveal sensor is a non-linear sensor of target position that has a central high-acuity (foveal) region surrounded by a lower-acuity region; this sensor is motivated by the density of photo-receptors in a biological visual system. The sensor is steerable, so the foveal region can be located where desired. We investigated both sensors in which the acuity of the foveal region was fixed and sensors in which the acuity of the foveal region could be adjusted. A narrow foveal region gives very accurate measurements of target position within a narrow range of positions, while a wide foveal region gives less accurate measurements within a wider range.

We have investigated the performance of a target tracking algorithm based on the foveal sensor. The non-linearity of the foveal sensor (and in particular the sensor gains that vary as a function of position) make application of a Kalman filter-based tracker very difficult. Prior to the Attentive Sensing project, significant effort had been devoted to developing data preprocessing strategies to match a Kalman filter to the foveal sensor. In the bulk of our work, we have used a Bayesian filter implemented by a particle filter (*i.e.* by sequential importance sampling). In a Bayesian filter, the estimator directly computes the posterior density of the system state given the observation sequence. The particle filter approximates the posterior density by weighted samples.

We have investigated the performance of a heuristic approach to foveal sensor configuration. In this heuristic, we center the foveal region on the predicted target position, and adjust the width of the foveal region to match the variance of the predicted position estimate.

The performance of a tracker that estimates a target's one-dimensional position and velocity using a foveal sensor was evaluated; we investigated both a fixed acuity sensor and an adaptive acuity sensor. As expected, the adaptive acuity sensor performed better than the fixed acuity sensors.

We have also investigated tracking a target moving in three dimensions using multiple foveal sensors. We have investigated tracker performance for both a centralized sensor configuration architecture and for a distributed sensor configuration architecture.

1.2 Personnel Supported

Personnel supported by this project include:

1. Darryl Morrell, PI

2. Doug Cochran, Co-PI

3. Dana Sinno, Graduate Research Assistant
4. Ya Xue, Graduate Research Assistant

1.3 Publications

The following conference papers were written as a consequence of this work:

1. D. Sinno, D. Cochran, and D. Morrell, "Multi-mode Detection with Markov Target Motion," *Proceedings of the Third International Conference on Information Fusion, Volume 2,* July 2000, pp. 25-31.
2. D. R. Morrell and Y. Xue, "Analysis of and Heuristics for Sensor Configuration in a Simple Target Localization Problem," *Conference Record of the 35th Asilomar Conference on Signals, Systems, and Computers,* November 2001, pp. 1391–1395.
3. D. R. Morrell and Y. Xue, "Bayesian Analysis of Target Localization using a Dual-Mode Sensor," *Proceedings of the 2002 International Conference on Acoustics, Speech, and Signal Processing,* May 2002, pp. 1593-1596.
4. Y. Xue and D. Morrell, "Adaptive Foveal Sensor for Target Tracking," *Conference Record of the 36th Asilomar Conference on Signals, Systems, and Computers,* November 2002, pp. 848–852.
5. Y. Xue and D. Morrell, "Target Tracking and Data Fusion using Multiple Adaptive Foveal Sensors," to be presented at the 6th International Conference on Information Fusion, July 2003.

The following theses and dissertations were supported in whole or in part by this grant:

1. Dana Sino, *Attentive management of configurable sensor systems*, PhD Dissertation, 2000.
2. Ya Xue, *Configuration of Systems of Attentive Sensors*, MS Thesis, December 2002.

1.4 Technical Interactions

The following research interactions were related to this project:

1. D. Morrell presented a talk in the ASU Math Department's Cognition Seminar on this work in March of 2001.
2. D. Morrell and Y. Xue presented a poster at the Asilomar conference on this work in November of 2001.
3. D. Morrell presented a talk in the ASU Math Department's Cognition Seminar on particle filters and the foveal sensor on April 9, 2002.
4. D. Morrell presented a talk at Brigham Young University on particle filters and the foveal sensor on May 29, 2002.
5. D. Morrell and Y. Xue presented a poster at ICASSP02 on this work.
6. D. Morrell presented a talk in the ASU Stochastic Modeling Seminar on particle filters and the foveal sensor on October 10, 2002.
7. Results from this work formed the starting point for the DARPA ISP contract that was awarded in July of 2002. The PI has begun work with Raytheon Missile Systems Division, who is the prime contractor on the DARPA ISP effort, to develop useful models for sensor scheduling. The PI has also provided some help to Raytheon in the area of target tracking using particle filters.

2 Introduction

The technical work done in this project was organized into two broad areas. The first was attentive detection and localization of a target using a configurable sensor; this area dealt primarily with problems in which the unknowns and sensor configuration variables were discrete. The second area was attentive estimation of the state of a dynamic system; this area involved unknowns and sensor configuration variables that were continuous. We briefly describe the progress in each of these two areas.

Our work on target detection and localization has included the development of several Matlab and C++ programs to evaluate the performance of sensor configuration algorithms. We have implemented simulators to evaluate the performance of the sensor configuration strategies. We have also implemented programs to directly compute system performance for the detection and localization problems; the programs provide exact values of the probability of error for a fixed number of sensor uses, and provide close lower bounds for the expected number of sensor uses needed to achieve a given probability of error. The work on detection and localization has been documented in two conference papers [1, 2].

Our work on attentive estimation also included the development of Matlab simulators; these simulators were used to ascertain the performance of state estimators using foveal sensor measurements for targets moving in one and three dimensions. This work has been documented in two conference papers [3, 4].

3 Attentive Detection and Localization

As previously described, attentive detection, localization, and classification problems typically involve discrete-valued unknowns and sensor settings. In a generic problem, the following three steps are repeated either a fixed number of times or until a given performance measure (*e.g.* maximum allowed probability of error) is achieved:

1. Configure one or more sensors (perhaps after selection from a suite of possible sensors).

2. Collect one or more observations using the sensors.

3. Use the observations to update the posterior probability distribution of the desired information.

This generic problem is a particular type of sequential Bayesian decision problem, in which current decisions (sensor configuration) may depend on previous observations. Sequential Bayesian decision theory is an appropriate approach for discrete-time, discrete-valued problems. Approaches to sequential Bayesian decision problems include probabilistic decision networks and dynamic programming for hidden Markov processes. Probabilistic decision networks have been developed under several different names, including Bayesian networks, factor graphs, and influence diagrams.

We have investigated the following target detection and localization problem: a target, located in one of M cells, is detected and/or localized by repeated use of a configurable sensor. Both single mode and dual mode sensors were investigated. The single mode sensor chooses a single cell to interrogate; the dual mode sensor chooses to interrogate either all cells simultaneously (Mode A) or a single cell (Mode B). We investigated two optimality criteria:

1. Minimize the probability of error when the sensor is applied a fixed number of times.

2. Minimize the number of sensor uses needed to achieve a given probability of error.

We used probabilistic decision networks to derive and understand the structure of myopic and multi-step optimal decision rules for optimality Criterion 1. We also investigated both optimality criteria using dynamic programming; a quantization method was used to approximately compute the optimal stationary decision

rule for Criterion 2. For Criterion 1, the performance of the myopic optimal rule was compared to that of the multi-step optimal rule and to that of several heuristics using both simulation and exact computation. For Criterion 2, the performance of several heuristics was compared. Some of the results of this work for the dual mode sensor are presented in Section 3.2.

In this section, we briefly review the tools of graphical models for decision processes and dynamic programming for hidden Markov processes. We then present results obtained for a dual mode sensor using these tools.

3.1 Tools to Implement Sequential Bayesian Decision Theory

We have used graphical models and dynamic programming (including the Partially Observed Markov Decision Process [POMDP] formalism) to solve and understand discrete-valued sequential sensor scheduling problems. For sequential decision problems of the type resulting from attentive sensor problems, the computational complexity of finding optimal rules typically grows exponentially, motivating the development and characterization of good heuristic configuration rules.

Graphical Models The use of graphical models to represent probabilistic relationships in inference problems has been developed in several different disciplines [5–8]. These models provide both a powerful method of representing complex probability distributions over many variables in terms of simpler relationships between a few variables, and a method to mechanize the computation of posterior distributions in response to observed variables [5]. Graphical probability models can, with the addition of decision nodes, implement optimal Bayesian decision rules [5, 9].

In the context of attentive detection, localization, and classification problems, the decision problem is specified in terms of the following model components:

1. Probabilistic models of unknowns. Unknowns are modeled using prior probability distributions; for example, prior target state information (*e.g.* presence or absence, position, ID, *etc.*) would be modeled by a prior distribution. Complex target and environmental models can be formulated using graphical methods.

2. Probabilistic models of sensor performance. The relationships between sensor settings, unknowns, and the resulting sensor observations are modeled by conditional probability distributions. These models may be derived from physical considerations or from statistical characterization of simulated or real sensors.

3. Cost and/or utility functions. Utility functions quantify the costs and benefits of decisions. For example, utility functions measure the benefits and costs of correct and erroneous decisions about the unknowns, the cost of different sensor configurations, *etc.*

Graphical methods provide tools for developing models and understanding their structure. Unfortunately, the solution techniques, while providing optimal solutions, do not typically discover computational reductions in the solutions that are available because of the structure of a specific problem. Additionally, the optimal solutions to problems involving multiple uses of a sensor typically grow exponentially in complexity as the number of sensor uses increases; graphical methods can make this problem clear, but may not provide insight how the structure of a given problem may be used to partially mitigate the exponential growth. Approximate computation in graphical models is an area of active research (*e.g.* [10, 11]).

Dynamic Programming for Hidden Markov Processes The repeated configuration and use of a sensor can be modeled as a hidden Markov process. The optimal sequence of sensor configurations can

be determined (at least in principle) using dynamic programming [12]. The conditional distribution of the unknowns given the information state (the entire past history of sensor configurations and acquired observations) is a sufficient statistic for the computation of the optimal sequence of sensor configurations; consequently, the optimal sensor configuration rule is a function of this conditional distribution. The dynamic programming solution typically involves computation of a cost-to-go function, which is computed recursively backward in time. The optimal sensor configuration strategy is usually a closed loop strategy, in which the optimal sensor setting at a given time depends on the observations obtained previously.

The methods used to compute the cost-to-go function and the associated optimal sensor configuration depend on the number of sensor uses. When the number of sensor uses is fixed before any observations are obtained, computing the optimal sequence of sensor allocations can be formulated as the solution to a Partially Observed Markov Decision Problem (POMDP) [13–15]. The solution of a POMDP is formulated in terms of value functions (which are essentially cost-to-go functions); under appropriate assumptions, the value function is piecewise linear and convex and can be computed as the solution of a linear program. Finding good approximate methods to compute value functions is a current area of research; see [16] for an overview of approximate methods.

When the sensor is repeatedly used until a stopping criterion is met, computing the optimal sensor configuration can be posed as an infinite horizon dynamic programming problem; solving this problem is more difficult than solving a POMDP. Numerical evidence indicates that the cost-to-go function for this problem is discontinuous, ruling out POMDP solution methods. We have implemented an approximate solution algorithm, in which the probability simplex is discretized to points on a lattice, and probabilities are quantized to these points. This discretization scheme is similar to that used in [17] and can be used to study the solutions to small problems.

3.2 Scheduling a Dual Mode Sensor

To illustrate the application of Bayesian tools to attentive detection and localization, we present the problem of scheduling a dual mode sensor. This work was completed as part of the Attentive Sensing project. Portions of this work were presented in [1,2].

In the dual mode sensor problem, a target may be present in an area to be searched by the sensor. The sensor is capable of operating in two modes: either interrogating the whole search area (Mode A) or interrogating a single cell in the search area (Mode B). In either mode, the sensor returns a binary-valued observation indicating whether the target was detected in the interrogated area. Our objective is to correctly detect and (if present) locate the target from a sequence of observations collected by the sensor.

The example presented here is an extension of the problem in [18–20]. In this previous problem, a general loss function was defined to reflect the desired outcomes of the sensor control problem. Also, the myopic optimal decision rule was derived and repeatedly applied to collect data until a stopping criterion was met. The performance of this approach was illustrated through several simulation examples, but no statistical characterization of performance over time was done. This work was also extended to include a Markov target motion model [21].

In this example, we use probability of error as our objective function. We also use a simpler model of sensor performance than that used in [19,20]; in our model, the sensor Mode B probabilities of detection and false alarm are the same for all cells. We derive the myopic optimal decision rule, and discover that it only depends on the probability that a target is present and on the probabilities of occupation of the two most probable cells. This is consistent with the results of [17,22], in which the myopic optimal rule was found for a single mode (Mode B) sensor; in [22], it was also shown that the myopic optimal rule is optimal over multiple sensor uses for the special case of a sensor with symmetric error performance.

Mathematical Model The target is either present and located in one of M possible cells or absent; the target state is a discrete random variable denoted $X \in \{0, \ldots, M\}$. A value of $X = 0$ indicates that the target is absent, while a value of $X = x \in \{1, \ldots, M\}$ indicates that the target is in cell x. Our knowledge of the initial target state is represented by a prior probability distribution $p_X(x)$.

The sensor output after the nth sensor use is the binary valued random variable Y_n. In Mode A, the sensor returns a value of $Y_n = 1$ if the target is detected in any cell and a value of $Y_n = 0$ otherwise. In Mode B, the sensor returns a value of $Y_n = 1$ if the target is detected in the interrogated cell and a value of $Y_n = 0$ otherwise. The choice of sensor mode and, in Mode B, the cell to interrogate is denoted $d_n \in \{0, \ldots, M\}$. A value of $d_n = 0$ indicates Mode A, while a value of $d_n > 0$ indicates Mode B interrogating cell d_n.

The sensor performance in each mode is known and characterized by a probability of correct detection and false alarm: p_{D_A} and p_{F_A} in Mode A, and p_{D_B} and p_{F_B} in Mode B. We define $\overline{p_{D_A}} = 1 - p_{D_A}$, $\overline{p_{F_A}} = 1 - p_{F_A}$, $\overline{p_{D_B}} = 1 - p_{D_B}$, and $\overline{p_{F_B}} = 1 - p_{F_B}$. The conditional distribution of Y_n given X and d_n is

$$p_{Y|X,d}(y_n|x, d_n) = \begin{cases} p_{D_A}, & y_n = 1 \text{ and } d_n = 0 \text{ and } x > 0 \\ \overline{p_{D_A}}, & y_n = 0 \text{ and } d_n = 0 \text{ and } x > 0 \\ p_{F_A}, & y_n = 1 \text{ and } d_n = 0 \text{ and } x = 0 \\ \overline{p_{F_A}}, & y_n = 0 \text{ and } d_n = 0 \text{ and } x = 0 \\ p_{D_B}, & y_n = 1 \text{ and } d_n > 0 \text{ and } x = d_n \\ \overline{p_{D_B}}, & y_n = 0 \text{ and } d_n > 0 \text{ and } x = d_n \\ p_{F_B}, & y_n = 1 \text{ and } d_n > 0 \text{ and } x \neq d_n \\ \overline{p_{D_B}}, & y_n = 0 \text{ and } d_n > 0 \text{ and } x \neq d_n \end{cases}$$

Each time the sensor is configured and an observed value y_n is obtained, y_n is used to update the distribution of X via Bayes theorem:

$$p(x|y_n, d_n, y_{n-1}, d_{n-1}, \ldots, y_1, d_1) = \frac{p_{Y|X,d}(y_n|x, d_n)\, p(x|y_{n-1}, d_{n-1}, \ldots, y_1, d_1)}{\sum_{x'} p_{Y|X,d}(y|x', d)\, p(x|y_{n-1}, d_{n-1}, \ldots, y_1, d_1)}$$

The sensor is used repeatedly either a fixed N times or until the posterior probability of a value x conditioned on all observations exceeds a threshold close to one. At this point, the estimated target location is the value whose posterior probability is largest; the estimated target location is denoted $\hat{x}$.

Optimal Rule for N Sensor Uses The optimal rule determines the sequence of decisions $d_1, \ldots, d_N$ that minimizes the expected probability of error. This rule can be easily derived using Bayesian decision networks [5] or dynamic programming [12]:

$$d_1, \ldots, d_N = \arg\max_{\tilde{d}_1} \sum_{y_1} \cdots \max_{\tilde{d}_N} \sum_{y_N} \max_{x} \left[\prod_{n=1}^{N} p_{Y|X,d}\left(y_n|x, \tilde{d}_n\right) \right] p_X(x) \tag{1}$$

The computation of this sequence is straightforward and can be implemented by directly implementing (1). Unfortunately, the computational complexity of this approach makes its application infeasible for all but small problems.

The myopic optimal rule (to minimize the probability of error after one observation) is a straight-forward extension of the myopic optimal decision rule for a single-mode sensor [17,22]. The structure of the problem allows significant simplification compared to a naive implementation of (1). The myopic optimal rule depends only on the probability that no target is present and the probabilities of the most probable and second most probable cells. The myopic decision rule is the following; for notational convenience, we drop the time subscript on d and the explicit conditioning of the distribution of X on past observations.

Let $\overline{m}$ and $\hat{m}$ be the cell with largest probability and the cell with second largest probability:

$$\overline{m} = \arg\max_{x>0} p_X(x)$$

$$\hat{m} = \arg\max_{\substack{x>0 \\ x\neq\overline{m}}} p_X(x)$$

Compute the following values:

$$L_0 = \max\left[\overline{p_{F_A}}\, p_X(0), \overline{p_{D_A}}\, p_X(\overline{m})\right] + \max\left[p_{F_A}\, p_X(0), p_{D_A}\, p_X(\overline{m})\right]$$

$$L_{\overline{m}} = \max\left[\overline{p_{F_B}}\, p_X(0), \overline{p_{D_B}}\, p_X(\overline{m}), \overline{p_{F_B}}\, p_X(\hat{m})\right] + \max\left[p_{F_B}\, p_X(0), p_{D_B}\, p_X(\overline{m}), p_{F_B}\, p_X(\hat{m})\right]$$

$$L_{\hat{m}} = \max\left[\overline{p_{F_B}}\, p_X(0), \overline{p_{F_B}}\, p_X(\overline{m}), \overline{p_{D_B}}\, p_X(\hat{m})\right] + \max\left[p_{F_B}\, p_X(0), p_{F_B}\, p_X(\overline{m}), p_{D_B}\, p_X(\hat{m})\right]$$

If $L_{\hat{m}}$ is largest, set $d_0 = \hat{m}$. Otherwise, if L_0 is largest, set $d_0 = 0$. Otherwise, set $d_0 = \overline{m}$.

Optimal Rule to Minimize the Average Number of Sensor Uses The optimal rule to minimize the number of sensor uses is an optimal stationary policy of an infinite horizon partially observed Markov decision process. Numerical simulations indicate that the optimal cost-to-go function is discontinuous (and hence not convex). Thus, the POMDP solution techniques based on piecewise linear functions cannot be used. We have found approximate optimal rules by quantizing the probability simplex into a lattice of points, and computing the optimal stationary policies and cost-to-go for the quantized simplex. The approximate optimal rules are functions of the probability distribution of X; based on these results, we believe that computationally tractable representations of these rules will be difficult, if not impossible, to find. This motivates the search for good suboptimal heuristics. Figure 1 shows a plot of the optimal decision function when $N = 2$ for the following sensor parameters: $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.95$, $p_{F_B} = 0.06$. The plot is a projection of the simplex onto the plane of the paper; the distribution $[p_X(0) = 1, p_X(1) = 0, p_X(2) = 0]$ is in the lower left corner, the distribution $[p_X(0) = 0, p_X(1) = 0, p_X(2) = 1]$ is in the lower right corner, and the distribution $[p_X(0) = 0, p_X(1) = 1, p_X(2) = 0]$ is in the top middle. Note that there are not easily defined decision boundaries, precluding a simple representation of the optimal policy.

Heuristics These heuristics are motivated by the myopic optimal solution (similar heuristics have been found to perform well for the simpler single mode problem in [1]).

1. Choose d_n as the most probable value of X.
2. Choose d_n as the the second most probable value of X.
3. Choose the myopic optimal control.
4. If the probability of target present is below a threshold close to one, choose $d_n = 0$; otherwise, choose d_n as the most probable value of X. This heuristic was used only when repeating observations until the probability of correct decision exceeds a threshold.

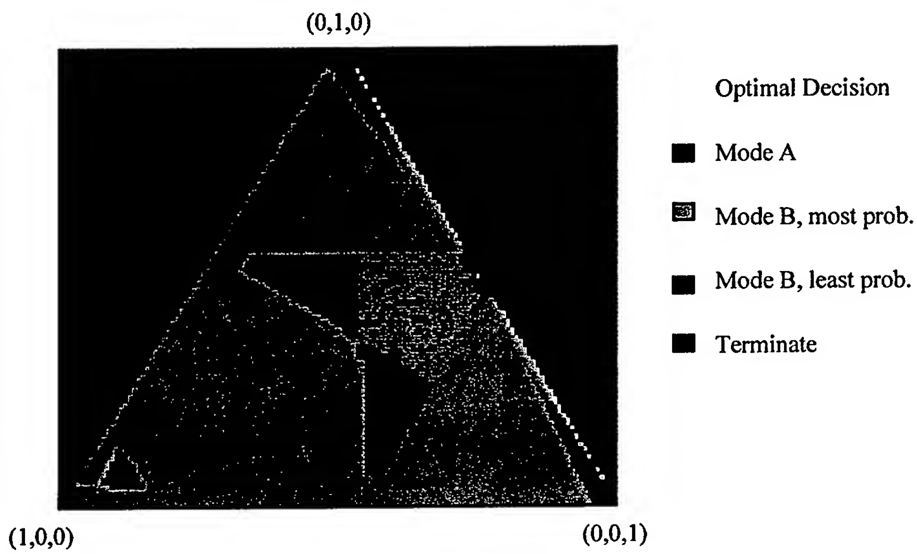

Figure 1: Optimal decision function: $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.95$, $p_{F_B} = 0.06$.

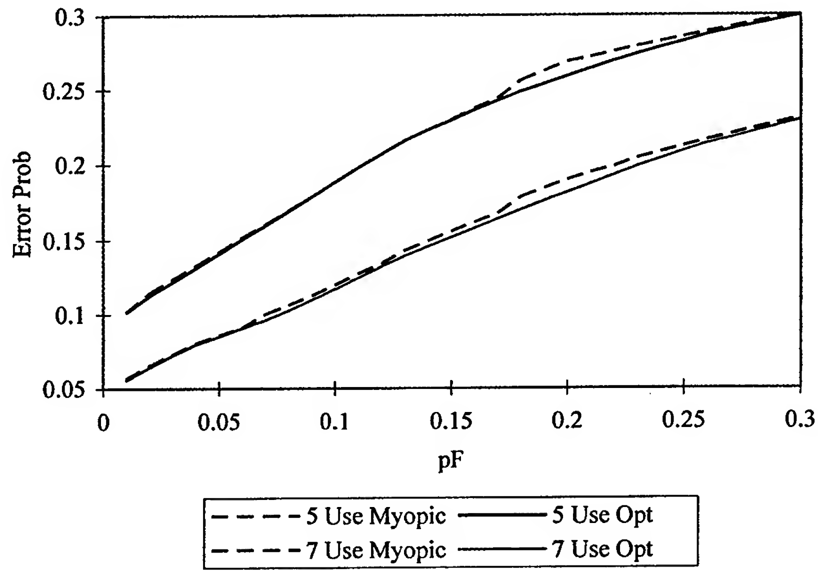

Figure 2: Comparison of the probabilities of error for the multi-step optimal and myopic optimal decision rules: $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.95$, p_{F_B} varied from 0.01 to 0.30. The number of cells was $M = 6$; the probability of error was evaluated for $N = 5$ and $N = 7$ sensor uses.

Computation of Performance The probability of error for N sensor uses given a particular control law is computed directly by considering a tree; each path from the root of the tree to a terminal node represents a possible sequence of observations. The probability of error can be computed by doing a depth-first traversal of the tree. The average number of sensor uses needed to guarantee a given probability of error is also computed by traversing a tree of possible observation sequences. In this case, the branches of the tree may have differing lengths (some branches may have infinite lengths, corresponding to observation sequences for which the decision rule does not terminate). In order to complete the computations in a reasonable amount of time, computations are truncated when the probability of the observation sequence associated with a branch in the tree drops below a specified threshold. Thus, the computed average number of sensor uses is lower than the actual expected value; we estimate that relative magnitude of this error is no more that a few percent.

Evaluation of Performance We have evaluated optimal and heuristic decision rules in terms of their probability of error and average number of observations needed. In particular, we have made the following evaluations:

1. Comparison of the probability of error of the multi-step optimal and the myopic optimal algorithms for a fixed number of sensor uses.

2. Comparison of the myopic optimal, most probable, and second most probable heuristics for a fixed number of sensor uses.

3. Comparison of all four heuristics to minimize the number of sensor uses.

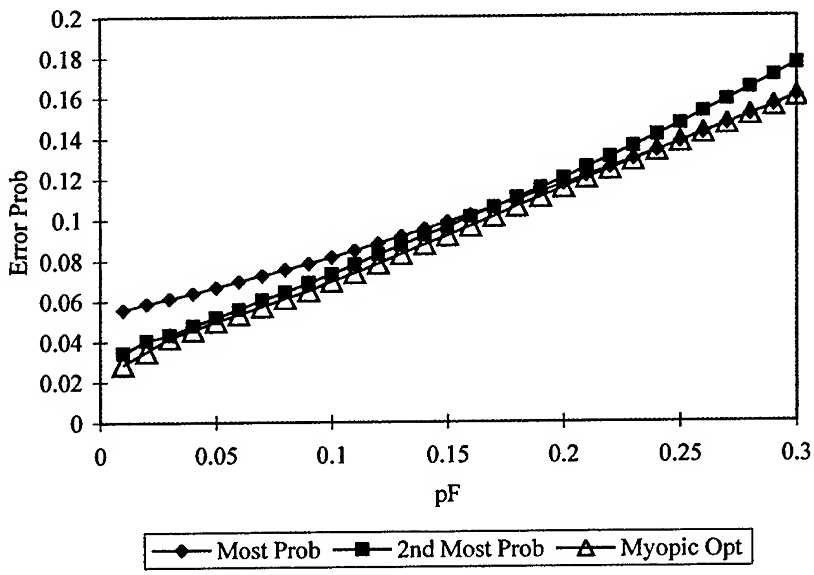

Figure 3: Comparison of the probabilities of error for the myopic optimal, most probable, and second most probable decision rules: $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.95$, p_{F_B} varied from 0.01 to 0.30. The number of cells was $M = 10$, and the number of sensor uses was $N = 15$.

4. Comparison of the most probable heuristic with the optimal rule to minimize the number of sensor uses.

The relative performance of the multi-step optimal and the myopic optimal algorithms for a fixed number N of sensor uses was computed. These computations were performed with $M = 6$ cells and values of N of 5 and 7. The probabilities of detection and false alarm for the two modes were set to $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.9$, and p_{F_B} was varied from 0.01 to 0.30. Figure 2 shows the probability of error as a function of p_{F_B}; note that the performance improvement of the multi-step optimal over the myopic optimal is not large.

The relative performance of the myopic optimal, most probable, and second most probable heuristics were also evaluated for a fixed number of sensor uses. This evaluation was performed with $M = 10$ cells and $N = 15$ sensor uses. The probabilities of detection and false alarm for the two modes were set to $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.9$, and p_{F_B} was varied from 0.01 to 0.30. Figure 3 shows the probability of error as a function of p_{F_B}. From the figure, it is clear that the myopic optimal decision rule has the best performance. For larger values of p_{F_B}, the performance of the most probable decision rule is the same as the myopic optimal because the myopic optimal almost always chooses the most probable value of X.

The expected number of sensor uses necessary to achieve a probability of error less than 0.05 was computed for the four heuristics. This computation was made with $M = 10$ cells, and the probabilities of detection and false alarm for the two modes were set to $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.9$, with p_{F_B} varied from 0.01 to 0.30. Figure 4 shows the average number of sensor uses as a function of p_{F_B}. Choosing the most probable value of X gave the best performance. The two-stage heuristic also gave good performance.

In order to characterize the sub-optimality of the heuristics in minimizing the expected number of observations necessary to achieve a given probability of false alarm, the cost-to-go functions for the optimal rule and for the most probable rule were computed using the probability discretization approach. Figure 5

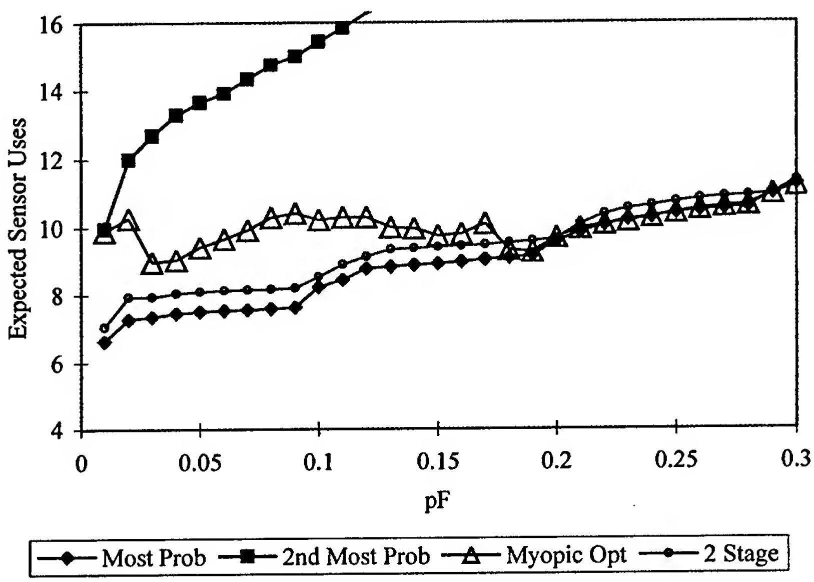

Figure 4: Comparison of the expected number of sensor uses for the myopic optimal, most probable, second most probable, and two stage decision rules: $p_{D_A} = 0.8$, $p_{F_A} = 0.15$, $p_{D_B} = 0.95$, p_{F_B} varied from 0.01 to 0.30. The number of cells was $M = 10$.

shows several histograms of the relative increase in the expected number of sensor uses for the most probable rule relative to the optimal rule over all points in the lattice. These results are for $M = 10$ cells. From these histograms, we see that the most probable heuristic is typically 20-30 percent worse than the optimal. Note that these values are computed using the quantized probability simplex and are approximations to the performance of the actual system.

Overall, these results show that heuristics with low computational requirements and good performance are often available. In particular, the myopic optimal solution appears to give good performance when minimizing the probability of error for a fixed number of observations; the myopic optimal heuristic has been used frequently [15, 17, 19–21]. Also, choosing the most probable value of X gives good performance when minimizing the number of sensor uses.

4 Attentive Estimation in Dynamical Systems

Discrete-time dynamical systems represent the dynamics of targets of interest. Attentive estimation of the state of these system models requires that sensors be configured before collecting observations. Thus, we have considered dynamic systems with configurable output maps; the sensor output maps may include non-linear relationships between observations and the system state. Our goal was to find control laws for sensor configurations to optimize state estimator performance.

We have used Bayesian filters implemented by particle filtering (sequential importance sampling methods) as state estimators. In a Bayesian filter, the estimator directly computes the posterior density of the system state given the observation sequence. For a given a cost function (*e.g.* squared error), the optimal state estimate can be computed from this density. Bayesian filters have many properties that are important for

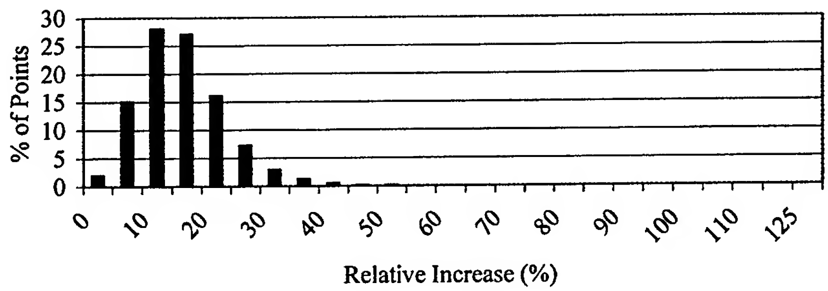

(a)

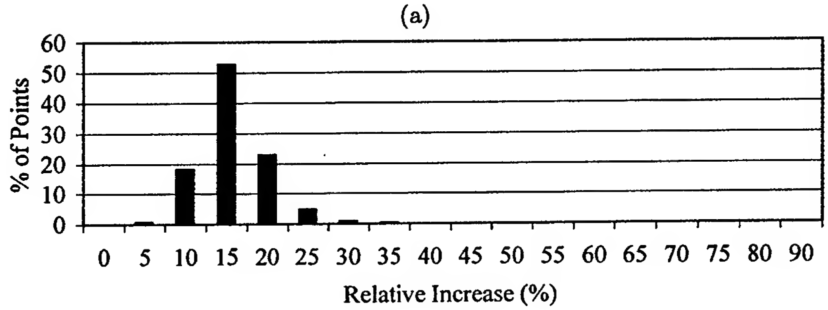

(b)

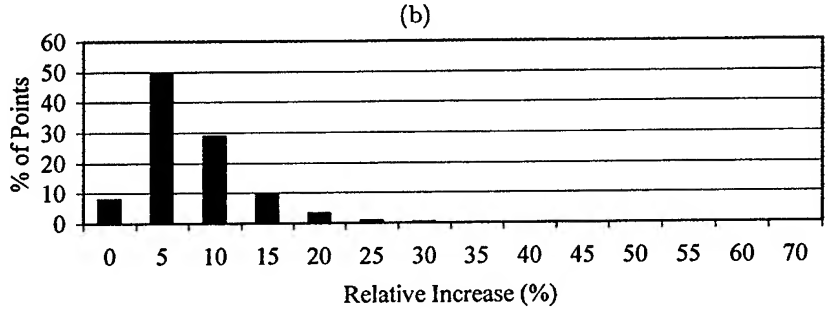

(c)

Figure 5: Histograms of the relative increase in the expected number of sensor uses for the most probable decision rule relative to the optimal decision rule. The number of cells was $M = 10$: (a) $p_{D_A} = 0.8$, $p_{F_A} = 0.2$, $p_{D_B} = 0.95$, $p_{F_B} = 0.00$, (b) $p_{D_A} = 0.8$, $p_{F_A} = 0.2$, $p_{D_B} = 0.74$, $p_{F_B} = 0.25$, (c) $p_{D_A} = 0.9$, $p_{F_A} = 0.1$, $p_{D_B} = 0.95$, $p_{F_B} = 0.06$.

the proposed work [23]: estimators for dynamic systems with non-linear state dynamics and observation maps can be formulated; the system state and the observation can include both continuous and discrete components; constraints on possible state values can be incorporated directly into the estimate; and multi-modal prior and posterior densities can be represented, so that observation maps that are many-to-one (*i.e.* that map many state values to a single observation) can be addressed in a straight-forward way. The primary difficulty in implementing a Bayesian filter is that closed form expressions for the posterior density typically do not exist. Thus, we use particle filtering (a form of Monte Carlo sampling) to implement these filters.

In this section, we briefly review foveal sensors and particle filters. We then present results obtained for a foveal sensor tracker.

4.1 Foveal Sensors

A foveal sensor has a high acuity region surrounded by a low acuity region. Foveal sensors are inspired by the characteristics of biological vision systems. Much of the work done in tracking targets with foveal sensors has been done in the context of computer vision (*e.g.* [24,25]), and most has been primarily concerned with the development of hardware systems that can track simple objects visually.

Theoretical work on tracking a target in one dimension with a foveal sensor is described in [26,27]. The emphasis of this work was the development of different approximation techniques to address the non-linear foveal sensor and, in the process, develop bounds on expected squared tracker error that could be computed using a Kalman filter. Most of the approximations were implemented using an observation data pre-processor followed by a Kalman filter. This work showed that the foveal sensor provided improved tracker performance (in the form of reduced squared error) over non-foveal sensors.

4.2 Particle Filtering

Particle filtering is a generic name for Monte Carlo sequential importance sampling methods to estimate the state of a dynamic system from noisy observations. Over the last decade, sequential importance sampling methods have been developed in many disciplines, including target tracking, image and video processing, signal processing, statistical mechanics, molecular biology, robot navigation and control, and artificial intelligence (reasoning under uncertainty). A comprehensive survey of the theory and application of particle filters is provided by [28], and an excellent tutorial introduction to state estimation using particle filters is provided in [29]. Particle filters and related Monte Carlo techniques are solving many target tracking and signal processing problems [30].

Let x_k denote the state of a discrete-time system at time k, and let z_k denote the observation of the system at time k. The discrete-time system model consists of a dynamic equation and an observation equation:

$$x_{k+1} = f_k\left(x_k, w_k\right)$$

$$z_k = h_k\left(x_k, v_k\right)$$

In this model, $f_k(\cdot)$ and $h_k(\cdot)$ may be non-linear functions of their arguments, and the noise values w_k and v_k may be non-Gaussian; for most target tracking problems, w_k and v_k are assumed to be white and independent of each other.

In a particle filter, the predicted and posterior densities $p\left(x_k|z_1,\ldots,z_{k-1}\right)$ and $p\left(x_k|z_1,\ldots,z_k\right)$ are approximated by a collection of pairs of samples and weights. The collection is denoted $\left\{x_k^{(i)}, w_k^{(i)}\right\}_{i=1}^{M}$, where $x_k^{(i)}$ denotes a sampled state value and $w_k^{(i)}$ denotes its importance weight. Each pair in this collection

is called a particle. The expected value of a function of the state is approximated by the following sum:

$$E[f(X_k)] \approx \sum_{i=1}^{M} w_k^{(i)} f\left(x_k^{(i)}\right)$$

An iteration of a particle filter consists of computing samples of the posterior density $p(x_k|z_1, \ldots, z_k)$ from samples of $p(x_{k-1}|z_1, \ldots, z_{k-1})$. This computation is typically accomplished by sampling from a proposal distribution, computing importance weights for the new sample, and then optionally resampling according to the importance weights. Typically, the proposal distribution is one from which samples can be drawn efficiently; $p\left(x_k|x_{k-1}^{(i)}\right)$ is often used as a proposal distribution. The importance weights depend on the actual value observed for z_k; when $p\left(x_k|x_{k-1}^{(i)}\right)$ is used as the proposal distribution, the importance weights are proportional to $p\left(z_k|x_k^{(i)}\right)$. Under reasonable conditions, it has been shown that this process converges asymptotically to the exact posterior distribution $p(x_k|z_1, \ldots, z_k)$ as the number of particles goes to infinity.

Several techniques have been developed to reduce the number of particles (and thus the computational load) needed to implement a particle filter with a desired error performance. One technique is Rao-Blackwellization, in which components of the state or observation that have closed form expressions for marginal densities are marginalized out. When the system state dynamics are linear with additive Gaussian noise, and the (non-linear) observation depends only on a component of the state, Rao-Blackwellization significantly reduces the number of particles needed in the computations. Also, the choice of the proposal distribution strongly affects the performance of the filter; one must often make a difficult trade off between the computational complexity of choosing samples from the proposal distribution and the number of particles needed to obtain a desired accuracy.

4.2.1 Example-Target Tracking Using a Foveal Sensor

We illustrate the potential benefits of target tracking with an attentive sensor by considering tracking of a target capable of 1-D motion using a foveal sensor. This is work completed under the Attentive Sensing project. Portions of this work were presented in [3].

Target Dynamics Model for 1-D Motion The target dynamics are modeled as a linear discrete-time system driven by white Gaussian noise. The system state at time t_k consists of two components and is denoted

$$\mathbf{x}_k = \begin{bmatrix} x_1(k) \\ x_2(k) \end{bmatrix}$$

The target dynamics equation is

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{g}w_k$$

where $\mathbf{A}$ is a 2×2 matrix, $\mathbf{g} = \begin{bmatrix} 0 & 1 \end{bmatrix}'$, and w_k is a scalar white noise process with variance Q. For an appropriate choice of $\mathbf{A}$, $x_1(k)$ models the target position on a line and $x_2(k)$ models the target velocity.

Foveal Sensor The foveal sensor is described by the following equation:

$$z_k = \tan^{-1}\left(C_k\left[x_1(k) - d_k\right]\right) + v_k$$

where z_k is the observation at time k, v_k is white Gaussian noise with covariance R, d_k is the location of the center of the foveal region, and C_k is the gain in the foveal region. Figure 6 illustrates the relationship between output and state for the foveal sensor.

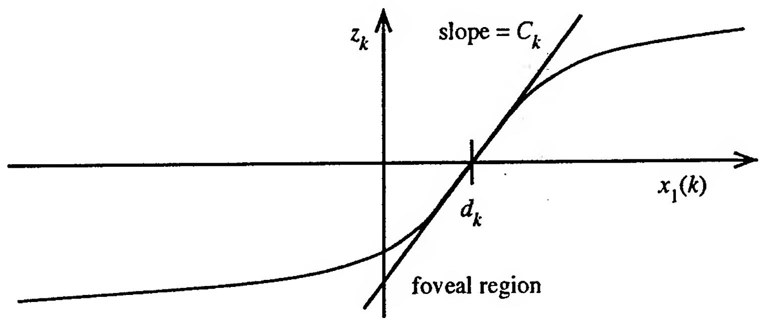

Figure 6: Output function of the foveal sensor; the foveal region is centered at d, and the acuity (gain) in the foveal region is approximately C_k, the slope of the output function at $x_1(k) = d_k$.

Configuring the Sensor Before obtaining an observation, the sensor is configured by choosing values for both d_k and C_k. We set d_k to $\hat{x}_1^-(k)$, the predicted estimate of $x_1(k)$ given the observations z_1 through z_{k-1}. (This value for d_k was also used in [26,27].)

Conceptually, the value of C_k is chosen so that a given percentage of the predicted particles will fall within the sensor's foveal region. In addition, the sequence of gains is smoothed using a simple exponentially weighted averaging filter. Thus, C_k is computed as

$$C_k = (1-\alpha)\tilde{C}_k + \alpha C_{k-1},$$

where $\tilde{C}_k$ is the instantaneous gain whose computation is described below, and α is the exponential weighting. By experimentation, we have determined that $\alpha = 0.8$ performs well.

The instantaneous gain $\tilde{C}_k$ is chosen after the value of d_k has been computed. Let w_k be the value for which a given percentage of the first component of the predicted particles falls into the interval $[d_k - w_k, d_k + w_k]$. Then the instantaneous gain $\tilde{C}_k$ is chosen so that the foveal region is $[d_k - w_k, d_k + w_k]$:

$$\tilde{C}_k = \frac{\pi}{2}\frac{1}{w_k}$$

For the results shown in Figure 8, w_k is chosen so that 80% of the predicted particles fall within the foveal region.

Performance Results for the 1-D Foveal Sensor We compared the particle filter implementation of the foveal sensor tracker with the best Kalman filter based tracker (of the several considered) in [27]. For this comparison, we used a system dynamic matrix

$$\mathbf{A} = \begin{bmatrix} 0.75 & 0.2 \\ -0.2 & 0.75 \end{bmatrix}$$

and a foveal sensor with both fixed and adaptive acuities. The gain of the fixed sensor was set to be consistent with the foveal sensor in [27].

The performance of the tracking algorithms was evaluated using Monte Carlo simulation; algorithms were compared on the basis of the average squared error in the tracker estimate of $x_1(k)$. Figure 7 shows curves of constant average estimation error as a function of the observation and driving noise variances. The particle filter fixed gain tracker performs as well as the best Kalman filter fixed gain tracker over most of the

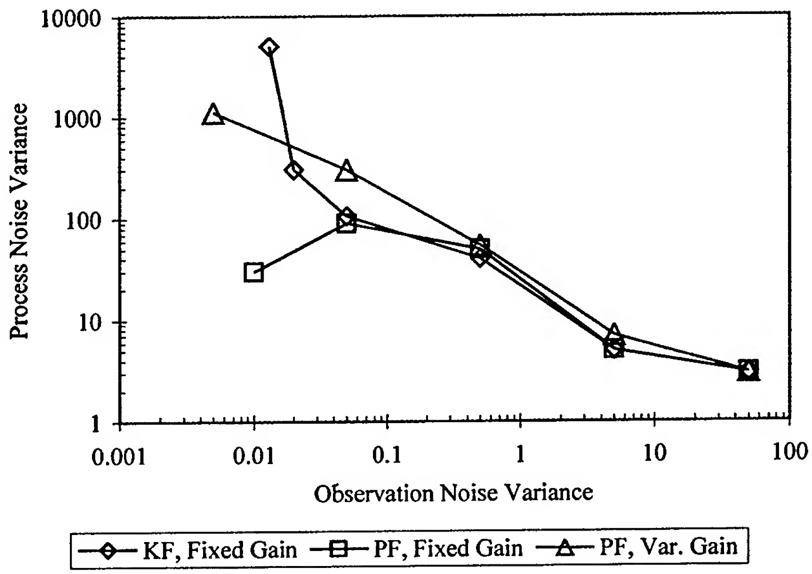

Figure 7: Curves of constant estimation error as a function of observation and process noise variances for the best Kalman Filter based tracker (KF, Fixed Gain), the adaptive acuity particle filter tracker (PF, Var. Gain), and the fixed acuity particle filter tracker (PF, Fixed Gain).

range of observation and process noise variances. The particle filter adaptive gain tracker performs better than the fixed gain trackers over most of the range of noise variances. The particle filter does not perform as well as the Kalman Filter when the observation noise variance is very small. This is because the observation density is sharply peaked for small variances, and few (if any) particles have weights significantly larger than zero. Several approaches have been developed to address this problem [28,29]; we are currently investigating whether they can be applied to the foveal sensor.

A more detailed comparison of the fixed acuity sensor and the variable acuity sensor was also conducted. The system dynamic matrix was

$$\mathbf{A} = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$$

These dynamics model one-dimensional motion with an interval of two seconds between observations. The process and observation noises were Gaussian white noise with variances $Q = 0.05$ and $R = 0.02$. The initial state was Gaussian with mean zero and variance one for both state components.

The average squared error in the estimate of $x_1(k)$ is shown in Figure 8 as a function of k for the adaptive acuity sensor and for three fixed acuity sensors. The fixed acuity sensors have gains of 0.25, 1.00, and 4.00. The squared error is averaged over 2,500 simulation runs. The adaptive acuity sensor performed better than any of the fixed acuity sensors. The tracker using the low-gain fixed-acuity sensor converges quickly but to a large average error. The tracker using the high-gain fixed acuity sensor converges slowly since the tracker searches for the target using a very narrow foveal region.

Extension to Three Dimensions We have extended the one-dimensional tracker to track targets in three dimensions. The three dimensional tracker uses three foveal sensors; each sensor measures azimuth

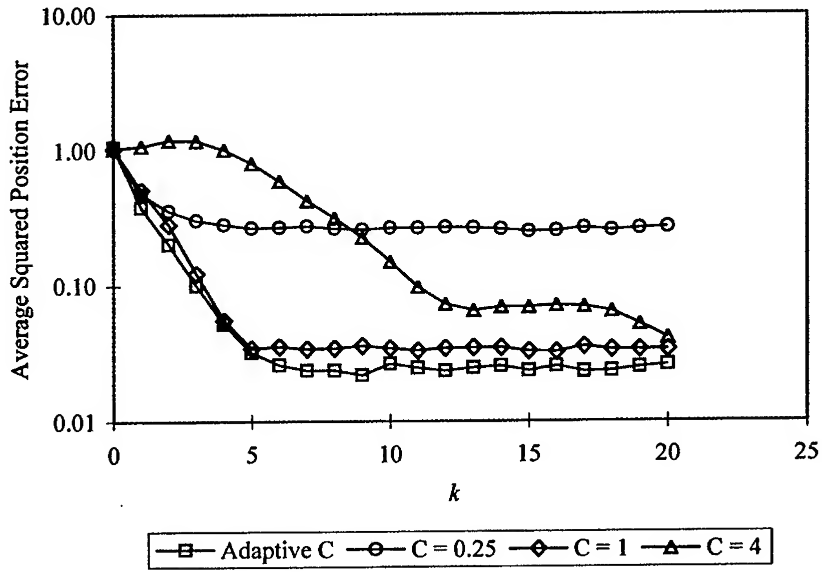

Figure 8: Average squared error in the estimate of $x_1(k)$ as a function of k for the adaptive gain and fixed gain sensors.

and elevation angles from the sensor to the target. Each sensor is aimed at the predicted position of the target, and the acuity of each sensor is adjusted individually using the algorithm described above.

We have investigated both global and local control schemes. In the global control scheme, the observations from all sensors are used by a global controller to estimate the target state; this state estimate is then used to configure all of the sensors. In the local control scheme, each sensor computes an estimate of the target state using only its observations; each sensor is configured using its local state estimate. Observations from each sensor are passed to a global estimator which estimates the target position. As expected, the global control scheme provides better estimates of the target state than the local control scheme. This improved performance comes at the cost of a significant increase in the information that must be communicated between the control node and the individual sensors.

References

[1] D. Morrell and Y. Xue, "Analysis of and heuristics for sensor configuration in a simple target localization problem," in *Conference Record of the 35th Asilomar Conference on Systems, Signals, and Computers*, pp. 1391–1395, Nov. 2001.

[2] D. Morrell and Y. Xue, "Bayesian analysis of target localization using a dual-mode sensor," in *Proceedings of the 2002 International Conference on Acoustics, Speech, and Signal Processing*, pp. 1593–1596, May 2002.

[3] Y. Xue and D. Morrell, "Adaptive foveal sensor for target tracking," in *Conference Record of the 36th Asilomar Conference on Systems, Signals, and Computers*, pp. 848–852, Nov. 2002.

[4] Y. Xue and D. Morrell. "Target Tracking and Data Fusion using Multiple Adaptive Foveal Sensors, "To be presented at Fusion 2003.

[5] R. Cowell, A. Dawid, S. Lauritzen, and D. Spiegelhalter, *Probabilistic Networks and Expert Systems.* Springer, 1999.

[6] B. Frey, *Graphical Models for Machine Learning and Digital Communication.* MIT Press, 1998.

[7] J. Pearl, *Probabilistic Inference in Intelligent Systems.* Mogan Kaufmann, 1988.

[8] R. Shachter, "Evaluating influence diagrams," *Operations Research*, vol. 34, pp. 871–882, 1986.

[9] S. Lauritzen and D. Nilsson, "Representing and solving decision problems with limited information," *Management Science*, vol. 47, pp. 1235–1251, September 2001.

[10] M. I. Jordan, Z. Ghahramani, T. Jaakkola, and L. K. Saul, "An introduction to variational methods for graphical models," *Machine Learning*, vol. 37, no. 2, pp. 183–233, 1999.

[11] M. Wainwright, T. Jaakkola, and Z. Willsky, "Tree-based reparameterization framework for analysis of sum-product and related algorithms," *IEEE Transactions on Information Theory*, vol. 49, pp. 1120–1147, May 2003.

[12] D. Bertsekas, *Dynamic programming: Deterministic and Stochastic Models.* Prentice-Hall, 1987.

[13] A. Cassandra, *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes.* PhD thesis, Brown University, 1998.

[14] M. Hauskrecht, *Planning and Control in Stochastic Domains with Imperfect Information.* PhD thesis, MIT, 1997.

[15] V. Krishnamurthy, "Algorithms for optimal scheduling and management of hidden Markov model sensors," *IEEE Trans. on Signal Processing*, vol. 50, pp. 1382–1397, June 2002.

[16] M. Hauskrecht, "Value-function approximations for partially observable markov decision processes," *Journal of Artificial Intelligence Research*, vol. 13, pp. 33–94, August 2000.

[17] V. Raghavan, M. Shakeri, and K. Pattipati, "Test sequencing algorithms with unreliable tests," *IEEE Trans. on Systems, Man, and Cybernetics-Part ?? A: Systems and Human*, vol. 29, pp. 347–357, July 1999.

[18] D. Cochran, D. Sinno, and A. Clausen, "Source detection and localization using a multi-mode detector: A Bayesian approach," in *Proceedings of ICASSP99*, vol. 3, pp. 1173–1176, Mar. 1999.

[19] D. Sinno, *Attentive Management of Configurable Sensor Systems.* PhD thesis, Arizona State University, May 2000.

[20] D. Sinno, D. Cochran, and D. Morrell, "A Bayesian risk approach to multi-mode detection," in *U.S./Australia Joint Workshop on Defense Signal Processing*, (Chicago), Aug. 1999.

[21] D. Sinno, D. Cochran, and D. Morrell, "Multi-mode detection with Markov target motion," in *Proceedings of the Third International Conference on Information Fusion, 2000*, pp. 25–31, July 2000.

[22] D. Castanon, "Optimal search strategies in dynamic hypothesis testing," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 25, pp. 1130–1138, July 1995.

[23] L. Stone, C. Barlow, and T. Corwin, *Bayesian Multiple Target Tracking.* Artech House, 1999.

[24] Y. Kuniyoshi, N. Kita, S. Rougeax, and T. Suehiro, "Active stereo vision system with foveated wide angle lenses," in *2nd Asian Conference on Conputer Vision*, pp. 359–363, 1995.

[25] D. Stack, C. Bandera, C. Wrigley, and B. Pain, "A real-time reconfigurable foveal target acquisition and tracking system," in *SPIE Proceedings: Acquisition, Tracking, and Pointing XIII*, vol. 3692, pp. 300–310, April 1999.

[26] D. Cochran and R. Martin, "Nonlinear filtering models of attentive vision," in *IEEE International Symposium on Circuits and Systems*, pp. 26–29, 1996.

[27] L. Li, D. Cochran, and R. Martin, "Target tracking with an attentive foveal sensor," in *Conference Record of the 34th Asilomar Conference on Systems, Signals, and Computers*, pp. 182–185, Nov. 2000.

[28] A. Doucet, N. De Freitas, and N. Gordon, eds., *Sequential Monte Carlo Methods in Practice.* Springer-Verlag, 2001.

[29] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, February 2002.

[30] "Special issue on Monte Carlo methods for statistical signal processing." *IEEE Transactions on Signal Processing*, February 2002.